Report on the Task Force on RHIC Detector Operations

On April 16, 2016 Associate Lab Director Berndt Mueller formed a task force to investigate the operational processes that were in place during the 2015 RHIC Run to ensure that the RHIC detectors were making optimal use of the resources devoted to their operation. The formation of the task force was occasioned by the recent realization that two recently installed subsystems, the PHENIX MPC-EX and STAR HFT, had major issues in their operational configurations during the entirety of RHIC Run 15, extending into the first few weeks of RHIC Run 16 in the case of the STAR HFT.   The task force was charged with identifying the root causes of the failures and recommending appropriate changes and improvements in the way the readiness and performance of detector systems at RHIC are monitored before and after RHIC runs in order to avoid similar failures in the future. The task force was composed of James Dunlop (chair), Michael Begel, Mickey Chiu, Bill Christie, Leo Greiner, John Lajoie, Laurence Littenberg, David Morrison, and Zhangbu Xu.

Root Causes

While the end effect was similar between STAR and PHENIX, the root causes of the data losses in Run 15 from STAR and PHENIX were different.

In the PHENIX case, it appears that the rush to commission a new detector while a run was ongoing led to the multiple failure modes observed in this type of environment:

- New and untested firmware and software was used for commissioning that needed to be revised as hardware problems were being worked through and the run was ongoing
- An uncalibrated detector at run start that made quantitative evaluations difficult
- An untested and unverified understanding of the detector element to electronics channel mapping generated confusion
- Unexpected and unplanned-for electronics problems with the SVX4 front end chips and associated RDO chain
- Unanticipated and unrelated issues from beam loss, diverting attention of the detector team
- Configuration parameters for the detector that were not recorded run by run during the commissioning period, making it difficult to retroactively track changes in detector performance

In the STAR case, a partial installation of the HFT had been commissioned in Run 13 and the full system operated in Run 14. A subtle RDO bug that led to an efficiency loss was introduced by a change in the RDO firmware between Run 14 and Run 15. This efficiency loss was noticed, but not attributed to the firmware change until a comparison between Run 14 and Run 16 results, under the same conditions, demonstrated that the problems associated with hit matching were not related to calibration or issues in the tracking algorithm. The RDO firmware is version controlled and archived, and tests were performed before placing changed firmware into production. These tests were inadequate to find the issue.

- The RDO firmware was extensively tested before use using both pattern data and the performance of full detector calibrations (threshold scans, noise rate measurements, etc.). None of these tests are sensitive to the firmware issue that was observed. The testing regimen in place was inadequate to find this problem.

Monitoring during the run

In both cases, the use of standard online monitoring software was ineffective in finding the problem. In the PHENIX case the hardware and firmware feeding the online monitoring was not yet in a state that allowed for conclusions to be reliably taken from the monitoring. Online monitoring could have isolated the problem if the full detector system had been in a more advanced state of preparedness. In the STAR case, online monitoring as it exists could not uncover this problem. Fast offline track matching was capable of finding this problem in the STAR context, through a comparison of hit matching between firmware versions for data taken under the same conditions, for example with cosmics. An efficiency loss in track to hit matching was seen in Run 15, but was incorrectly attributed to recoverable issues in calibration and tracking algorithms, and not successfully traced to the firmware change until Run 16.

Proposed changes

Appropriate changes and improvements in the way the readiness and performance of detector systems at RHIC is monitored before and during RHIC runs

- Time for detector commissioning with a sufficiently long prototype/commissioning period, up to and including an entire RHIC run, should be added appropriately to any new detector construction project. For this commissioning period, a significant portion of the new detector system should be in place, with its readout fully integrated into the experiment, along with a complete suite of monitoring processes.
- Online and offline software frameworks should be in place prior to detector commissioning, at minimum to the extent that they allow for meaningful testing of correlations with other detectors
- Firmware and software used in detector system hardware configuration and readout must be archived and versioned for all running configurations in a code

management system (CVS, SVN, etc.) Parameters used as input to the firmware must also be archived with a method appropriate for complete retrieval at a later date.

- All changes to the detector operation firmware/software need to be fully tested through a suite of regression tests in a full hardware/firmware/software environment that can be configured to demonstrate complete required functionality.

Detailed Description of the Operational Issues

STAR HFT

Problem description, how it was discovered and current running status.

A fast analysis of STAR Run16 Au+Au data taken near the beginning of the run showed a (tracks with at least 3 HFT hits / TPC tracks) matching ratio lower than the ratio observed in the Au+Au Run14 data set. The difference was significant, more than a factor of two (see figure below). This observation triggered a series of checks on the PXL detector configuration and data processing chain. By switching between different PXL firmware versions and taking data during the early part of the current RHIC Run in similar conditions of luminosity and trigger rate, it was found that the data taken with the firmware used during Run15 and in the first 3 weeks of Run16 (revision r865) displays a significant drop in the tracking efficiency with respect to the data taken with the firmware used during Run14 (revision r700).

The loss of efficiency became visible as soon as the first Au+Au data from Run16 were reconstructed, analyzed and compared to Run 14 results. In Run 15, the low efficiency in p+p and p+Au data was noted but not immediately attributed to effects resulting from the change in the firmware. Part of it was due to the different colliding systems and lack of reference and experience with the low multiplicity, high pile-up rate p-p environment, but also due to the different (and evolving) tracking software used for the data reconstruction, and the preliminary status of the STAR detector calibrations and alignment. *The low matching ratio was attributed to tracking software inefficiencies.* Full detector simulations were not useful at the time since already in Au+Au system a discrepancy between data and simulations was present and under investigation. Those discrepancies were resolved (decoder issue) early in run 16, after the Run15 data taking. At the same time the Run15 calibrations were performed and a detailed look was possible.

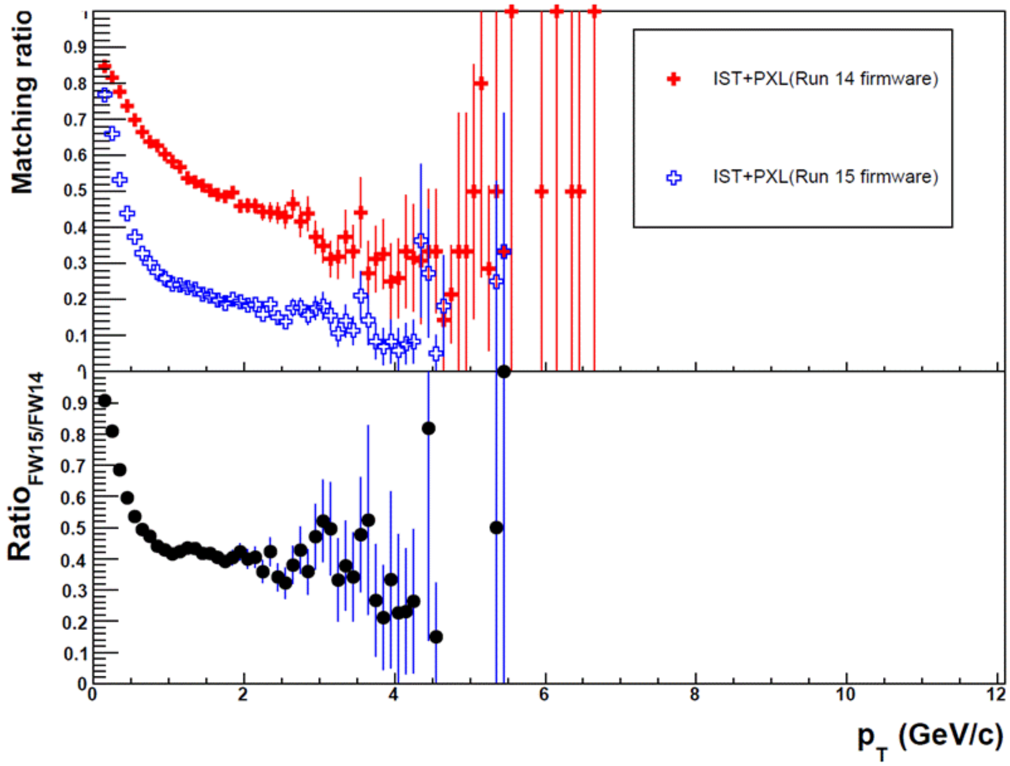

FIGURE 1: HFT/TPC TRACK MATCHING RATIO AS A FUNCTION OF THE TRACK MOMENTUM. ALL PLOTS MADE WITH RUN 16 DATA. THE RED SOLID CROSSES REFER TO THE DATA TAKEN WITH THE FIRMWARE USED IN RUN 14 (R700), BLUE OPEN CROSSES TO THE RUN 15 FIRMWARE (R865). THE RUN 15 FIRMWARE MATCHING EFFICIENCY NORMALIZED TO THE RUN 14 FIRMWARE EFFICIENCY IS SHOWN IN THE BOTTOM PANEL.

When spotted, as remediation to the observed effect, STAR immediately rolled back to Run14 firmware (r700) for the rest of Run16 data taking. At our current state of knowledge, this is the most efficient and safest PXL configuration and allows for a track matching efficiency as large as the one observed in Run14. Recent simulations reproduce the current PXL performance within 5%.

PXL Detector Characteristics and RDO

The PXL detector uses Monolithic Active Pixel Sensors to detect charged particle hits. Charged particle tracks leave a cluster of hit pixels (approximately 2.5 pixels/cluster) in the detector and the hit locations are read out as digital address data directly from the silicon sensors. The sensors in use at STAR have an integration time of 186 μs. This means that there is significant pile up in the detector which is resolved to hits of interest by tracking inward from the outer detectors with graded resolution until one is able to exclude most of the pileup and noise hits from the search window on each of the PXL layers. Since the data from the sensors is divided into frames (one full pixel array of hit data) and the frame boundaries are not aligned to the random STAR triggers it is the task of the RDO firmware to select the proper address ranges of hits that are associated with the trigger. This firmware task is accomplished in the FPGA by combining the knowledge of the frame phase (which row of the sensor is being digitized at the time) of the sensors on a sector and the length of the integration time. When a trigger is received the RDO firmware sets a pointer corresponding to the current internal address position of the digitation in the sensors, opens a memory buffer, and saves only the addresses that correspond to one integration time from that point. These hit addresses are then packaged into an event and sent off to STAR DAQ.

Processes in place to monitor the detector status during data taking

For all runs, there is a set of real time online monitoring processes that show the status of the detector. These consist of updating histograms and scatter plots that show the hit densities of each PXL sensor, summed over ladders and as a function of theta and phi. We also monitor and reported errors in the sensors and RDO system (overflow, loss of communication lock, etc.) and the dead time. Tracking in the High Level Trigger (HLT) in STAR does not include HFT at this time. Track matching between the TPC and HFT requires offline calibration and analysis. The calibration procedure shown below is done in a fast offline mode where the calibrated HFT hit data is matched to TPC tracks in a separate analysis process. The problem that was discovered is very subtle and not visible in any online analyses. The online processes display event by event and integrated hit densities and error conditions. These observables were the same for both firmware versions. In the case of what was discovered, it appears likely that valid hits in time are being lost but only partially as there are still valid matches. This is a very unusual type of error.

Run 15

During the Run 15 data taking in STAR the online monitoring processes were in place and working. Offline calibration, however, took a bit longer than expected (see below). Different beams required that different tracking software be used for this run in order to optimize for the lower multiplicity in the p+p and p+Au data. When the HFT calibration was accomplished, an inefficiency was noted in the cosmic ray tracking and in spot checks of the track matching with the TPC, but since this run involved new and un-tuned software, the inefficiency was initially attributed to the new tracking software or to the insufficient calibration of the TPC or HFT. This delayed the discovery of the problem.

Run 16

In run 16, the calibration procedure was accomplished in the first few weeks of running. The Run 16 data were directly comparable to the Run 14 data (with top energy Au+Au beams) and the already tuned high multiplicity tracking software was used for fast offline analysis. The Run16 to Run 14 comparison revealed the track matching efficiency loss. At that point an analysis of the cause was initiated. A report on the investigation of the loss in efficiency can be found at: https://drive.google.com/file/d/0BxS_WKusxUlxc05VY3JtNl9wWWs/view?usp=sharing. In Run 16, the fast offline analysis was used extensively to compare different RDO firmware versions and was a very useful tool in the investigation and remediation of the observed inefficiency. It is important to point out that the fast offline analysis tool used in Run 16 is a direct product of the software developed for the analysis of Run

14. Run 15 did not benefit from having software already prepared and tuned and STAR was developing this software at the time.

Calibration Procedure for Track Matching

The calibration of the PXL detector is a lengthy process. Preparing the calibration HFT for track matching has three primary steps. Since the PXL detector accumulates some damage during runs, we repair/replace ladders on the sectors between the runs. The positions of the sensors/pixels on the refurbished sectors are measured with a Coordinate Measuring Machine (CMM) and this data is put into the geometry description files. The PXL halves are physically removed and inserted after and before each subsequent run. At the beginning of each run period, the relative positions of the detector halves needs to be measured and calibrated with cosmic rays. Thus, the procedure to generate an internally understood and calibrated detector that can be used for proper track matching includes:

- Take noise runs with the newly refurbished detector to both generate the sensor threshold settings and to find hot and dead pixels/rows/columns. This data is put into masking tables that are applied to the raw hit data to give the hit data that will be used to perform clustering and generate good hit addresses.
- Take data with cosmic rays and perform a residual minimizing fit to cosmic ray tracks that aligns the positions of the detector sectors.
- Tune the tracking software parameters to maximize the efficiency for track matching.

It is also important to note that the TPC must be calibrated as well since it will have an impact on observed HFT matching ratio.

Short Discussion on Causes

The delay in the discovery of the track matching inefficiency seen in Run 15 and the beginning of Run 16 at STAR appears to be the result of several factors.
The effect of firmware and other changes to the system were not adequately checked through a full hardware/firmware simulation chain that allowed for changing event by event data in time relation to applied triggers. The existing checks of the readout chain included fixed patterns that are fed through the system, but in this case this testing was insufficient to uncover the problems that occurred. A full hardware/firmware simulation chain with the requisite properties is in development.
The delay in the discovery of the detector inefficiency appears to reflect on the subtlety of the problem discovered, the lack of software tuned for the p+p environment which was under development at the time and the availability of personnel trained in the TPC/HFT tracking software and HFT calibration procedures.

PHENIX MPC-EX

During Run-15, the MPC-EX timing was initially adjusted to the same Level-1 delay as the VTX strip-pixel detectors (40 BCLKs), as both systems utilize the SVX4 chip for digitization and readout of their respective Si detectors. By referencing the MPC-EX to the nominal VTX timing, it was believed that a scan of a few BCLKS around the nominal position should identify the correct beam crossing. The initial timing of the detector was done using p+p collisions at 200 GeV in early March, 2015, and due to the relatively low occupancy of the MPC-EX in p+p collisions the timing was evaluated by looking at the MPC-EX ADC spectra in each layer to identify the MIP peak. The MIP peak location itself was well-known based on studies with cosmics during detector assembly. Initial timing scans clearly identified a crossing within the expected range for both the north and south MPC-EX that showed a clear MIP peak in the data, and the MIP peak yield degraded rapidly when the timing was moved away from this crossing.

Based on the initial timing study attempts were made to correlate the number of hits in the MPC-EX with the BBC charge and to correlate clusters of hits with MPC showers. During Run-15 these correlation attempts were complicated by several key issues:

- The detector was completely uncalibrated, and methods for identifying hot channels and calibrating the relative energy response of the minipads were being developed on the fly. For this reason, correlations with the BBC were ambiguous as such correlations were as broad along the diagonal as they were perpendicular to the diagonal. In this case the low multiplicity in p+p made it very difficult to determine the correlation with an uncalibrated detector.
- There were serious readout and firmware issues affecting ~27% of the detector. The main problem was a "Cell ID lockup" in SVX4's at specific locations along the readout chain. Methods were still being developed to eliminate data from chips that were locked up, complicating the analysis. (This issue has been addressed for Run-16.)
- The detector geometry was still uncertain at the time. This, combined with the dead areas in the detector due to the Cell ID lockup issue described above made correlations with the MPC broader than expected in simulations, and it was difficult to determine if this was an artifact of geometry and readout issues or a result of random matches between the MPC-EX and MPC.
- Beam abort issues occupied a significant fraction of the time available to the MPC-EX group, as they struggled to evaluate the effect of each beam abort on the MPC-EX. After every beam abort a significant effort had to be made to determine if any damage was done to the detector, which distracted effort from evaluating the physics quality of the data being taken.

Based on the available analysis of the detector at the time in Run-15, with preliminary software and calibrations, it was concluded by the MPC-Ex group that the detector was properly timed in.

Between Run-15 and Run-16, substantial progress was made in detector calibration, a new offline software framework, simulations, and improved firmware for the detector. The new firmware, in particular, included features that made identification of the actual latency from a fixed signal (the reset of the SVX4 pre-amplifiers, triggered by a mode bit) much more clear. This new feature made it clear that the detector was actually operating at a latency of 20BCLKs, instead of 40BCLKs as expected.

Subsequent investigation of the SVX4 configuration, a serial bit stream that is downloaded to each chip, showed that the bit string had been incorrectly edited when programming the latency, and was shifted by a single bit, resulting in the desired latency of 40BCLKs being programmed as 20BCLKs. The download of the SVX4's is logged by the PHENIX "feed" program, and investigation of these logs led to the conclusion that the serial string was set this way since mid-March 2015, just after the initial timing runs. (The timing runs themselves were done with preliminary software and firmware and the logging capability had not yet been instituted. A second set of timing runs in late March, 2015 to double-check with the production firmware did use the incorrect SVX4 latency setting.)

Upon this discovery, a set of investigations were launched to look into the Run-15 data and verify the suspected mistiming. The plots below document the result of this investigation. A reference run (428385) from early in the p+p running period was chosen as a starting point for these investigations. This run has been fully calibrated in the MPC-EX offline framework, and therefore provides the best opportunity to look at both simple quantities (like minipad energies) as well as more processed quantities (like the number of EM showers found by the MPC-EX reconstruction).

Figure 2 shows the distribution of events with at least one MPC-EX shower found by the offline reconstruction software as a function of the GL1 crossing number. The mistiming of the detector can be seen as a "reflection" of the abort gap in this distribution, approximately 20BCLKs earlier than the true abort gap.

Figure 3 shows the MPC-EX minipad energy distribution (vertical axis, in arbitrary units) as a function of the GL1 crossing number. This is a less "processed" value than the number of showers found by the reconstruction software, and it clearly shows a depletion on the energy distribution in a region 20BCLKs prior to the abort gap, indicating the loss of showers in Figure 1 is not an issue with the shower reconstruction software but is seen in the energy measured in the detector.

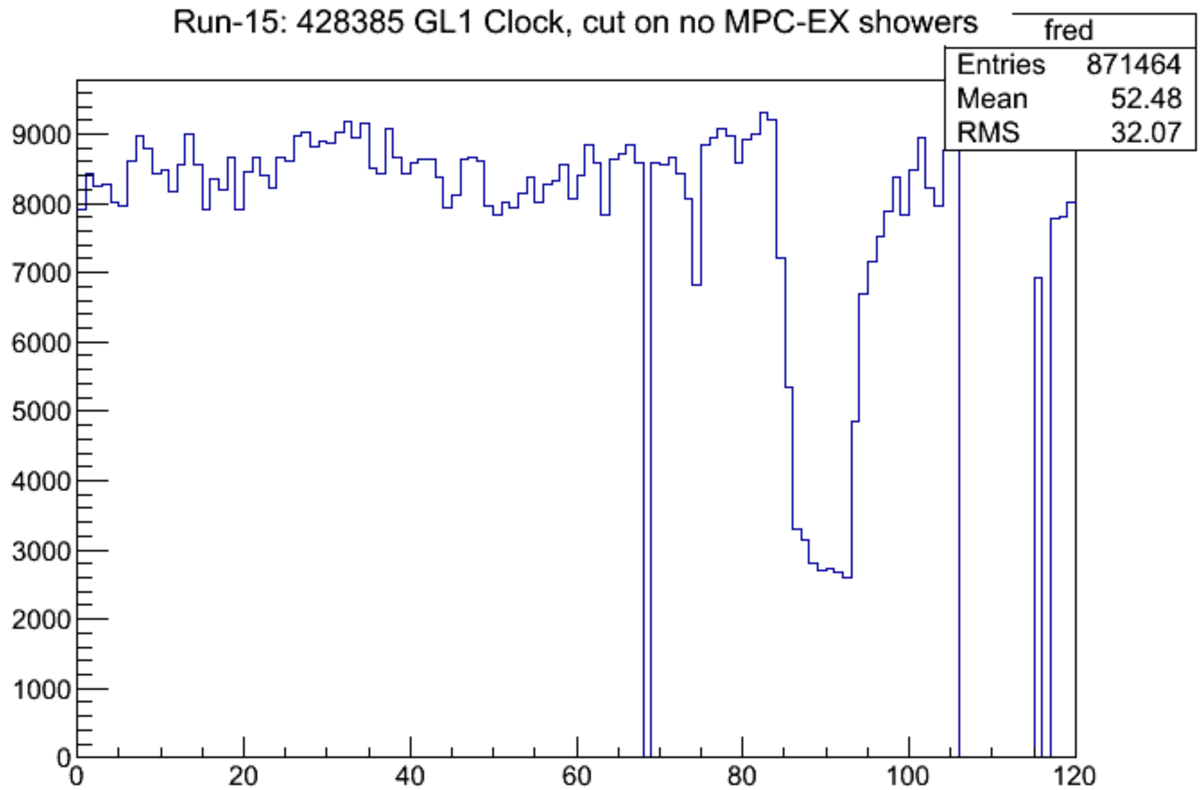

FIGURE 2: THIS IS THE NUMBER OF EVENTS WITH AT LEAST ONE MPC-EX SHOWER AS A FUNCTION OF THE GL1 CROSSING NUMBER. THE ABORT GAP IS CLEARLY VISIBLE STARTING AT CROSSING 107. THIS REGION IS DEPOPULATED BECAUSE THERE ARE NO TRIGGERS IN THIS REGION. THE REDUCTION OF SHOWERS ROUGHLY 20 BCLKS EARLIER IS THE REFLECTION OF THE ABORT GAP AS SEEN BY THE MPC-EX TIMING, AND THE REDUCTION OF SHOWERS IS DUE TO THE LACK OF COLLISIONS IN THE ABORT GAP. THIS IS STRONG EVIDENCE FOR THE MISTIMING OF THE MPC-EX.

In order to examine the correlation with the BBC the fully calibrated MPC-EX energy sum is compared with the BBC charge for the south arm of the PHENIX detector (Figure 4). The MPC-EX calibrations include the elimination of hot channels and hits from SVX4 chips with CellID problems. This correlation, based on calibrated data, shows that the BBC and MPC-EX are not correlated – they are not looking at the same beam crossing. Note, however, the extent of the correlation in p+p, as compared to Au+Au (see Figure 6).

In Figure 5 we show a spatial correlation between MPC-EX showers and MPC clusters in real data and in simulations. As with the BBC correlations, these results support the idea that the MPC-EX is mistimed with respect to the MPC as the correlations are consistent with random association.

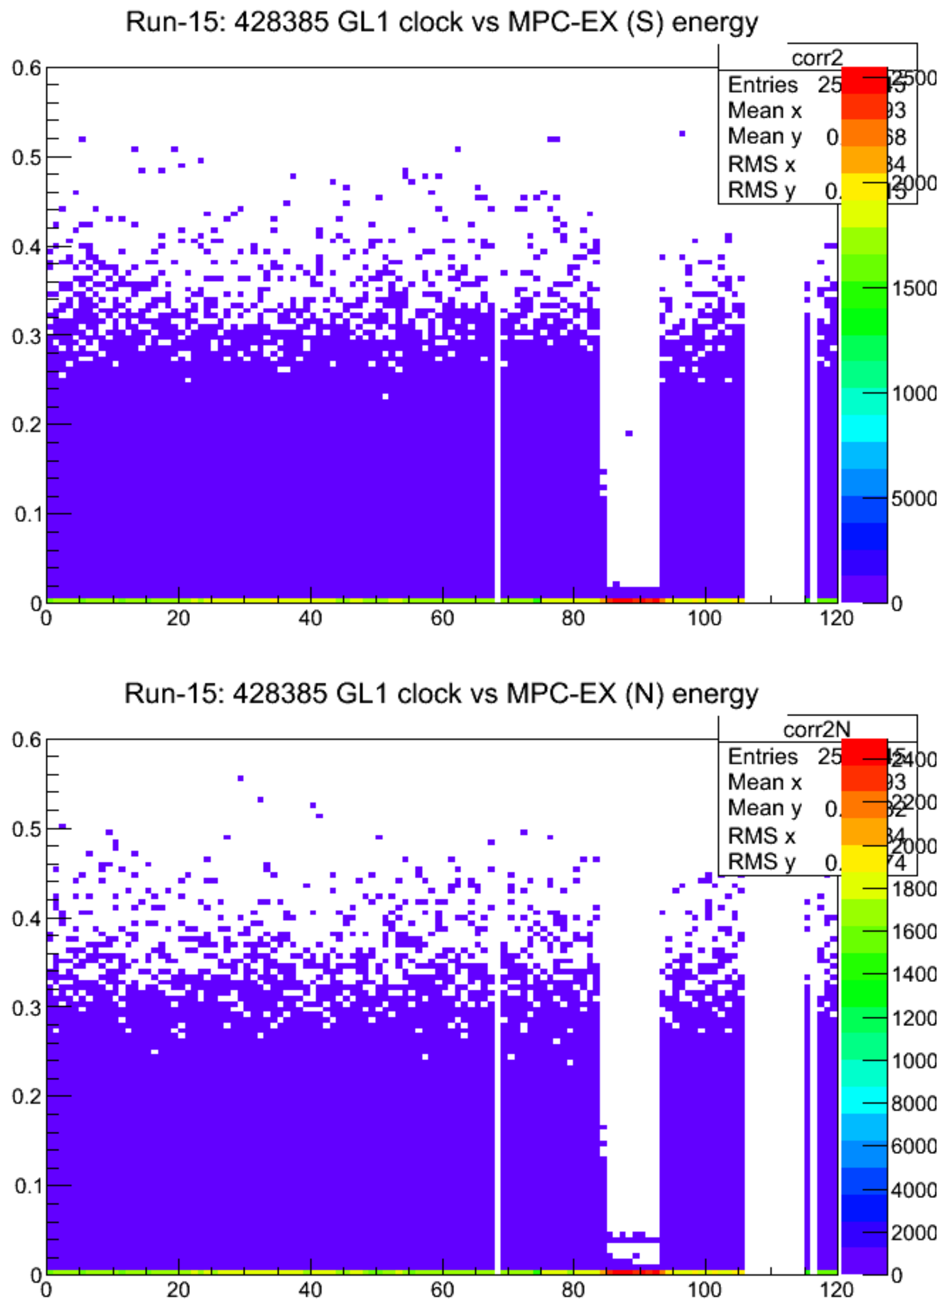

FIGURE 3: THIS FIGURE SHOWS THE MPC-EX MINIPAD ENERGY DISTRIBUTION (VERTICAL AXIS, IN ARBITRARY UNITS) AS A FUNCTION OF THE GL1 CROSSING NUMBER. THIS IS A LESS "PROCESSED" VALUE THAN THE NUMBER OF SHOWERS FOUND BY THE RECONSTRUCTION SOFTWARE, AND IT CLEARLY SHOWS A DEPLETION ON THE ENERGY DISTRIBUTION IN A REGION 20BCLKS PRIOR TO THE ABORT GAP, INDICATING THE LOSS OF SHOWERS IN FIGURE 1 IS NOT AN ISSUE WITH THE SHOWER RECONSTRUCTION SOFTWARE BUT IS SEEN IN THE ENERGY MEASURED IN THE DETECTOR.

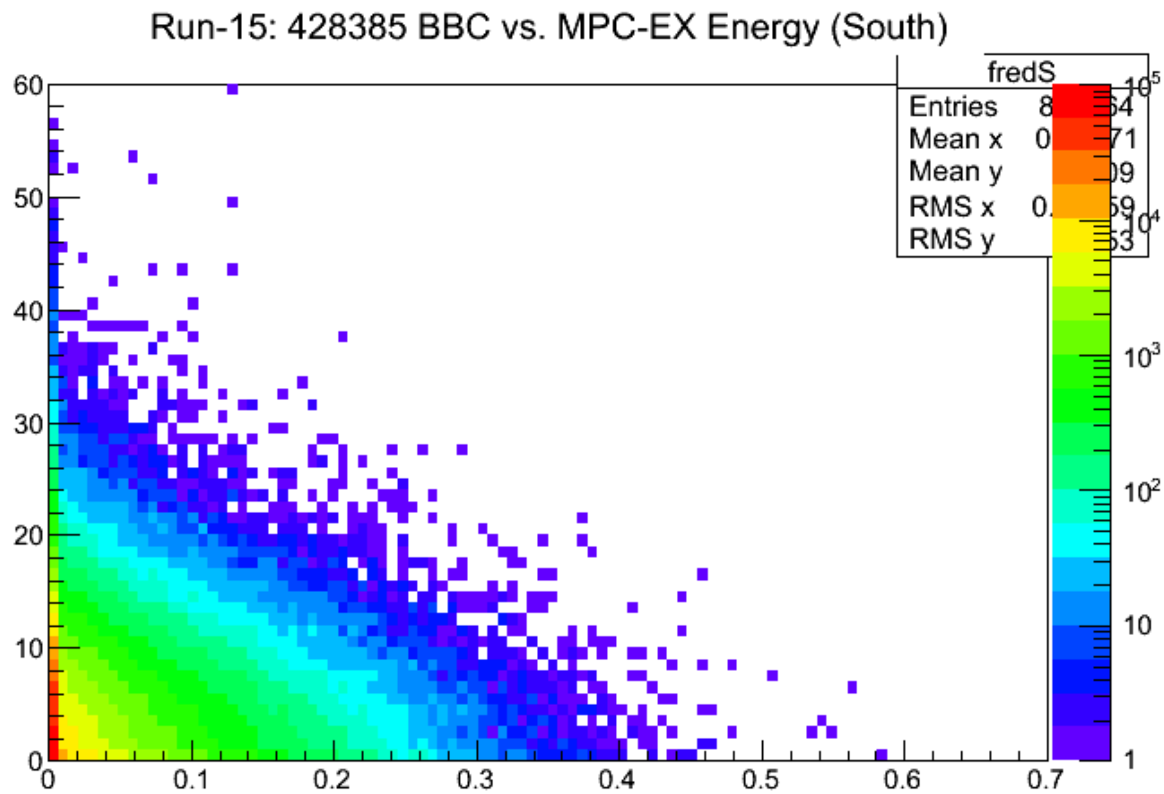

FIGURE 4: THIS FIGURE SHOWS THE CORRELATION BETWEEN THE BBC CHARGE (VERTICAL AXIS) VERSUS THE MPC-EX SUMMED MINIPAD ENERGY (HORIZONTAL AXIS) FOR THE SOUTH ARM. THIS PLOT IS A RESULT OF THE NEW OFFLINE FRAMEWORK WITH FULL CALIBRATIONS, AND HOT AND DEAD CHANNELS REMOVED.

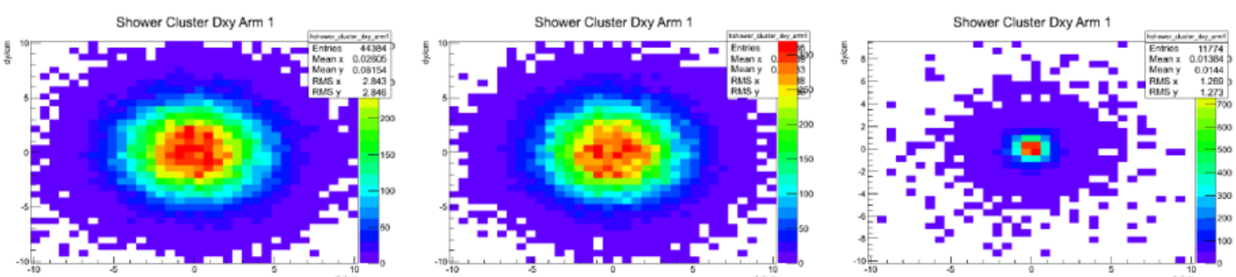

MPCEX pairs to the previous MPC event   MPCEX pairs to the same MPC event   Simulated Matched Event

FIGURE 5: A COMPARISON OF THE SPATIAL MATCHING IN X (HORIZONTAL) AND Y (VERTICAL) BETWEEN AN MPC-EX SHOWER AND A CLUSTER IN THE MPC. THE CENTER PANEL SHOWS MATCHING IN RUN-15 DATA IN THE SAME EVENT, THE LEFTMOST PANEL SHOWS MATCHING TO THE PREVIOUS EVENT (DELIBERATELY MISTIMED) AND THE RIGHTMOST PANEL SHOWS THE EXPECTATION IN SIMULATIONS, INCLUDING THE MPC-EX DEAD AREAS IN RUN-15. THE CONCLUSION IS THAT IN RUN-15 DATA THE MATCHING IS CONSISTENT WITH RANDOM ASSOCIATION.

Run-16 MPC-EX Timing Exercise

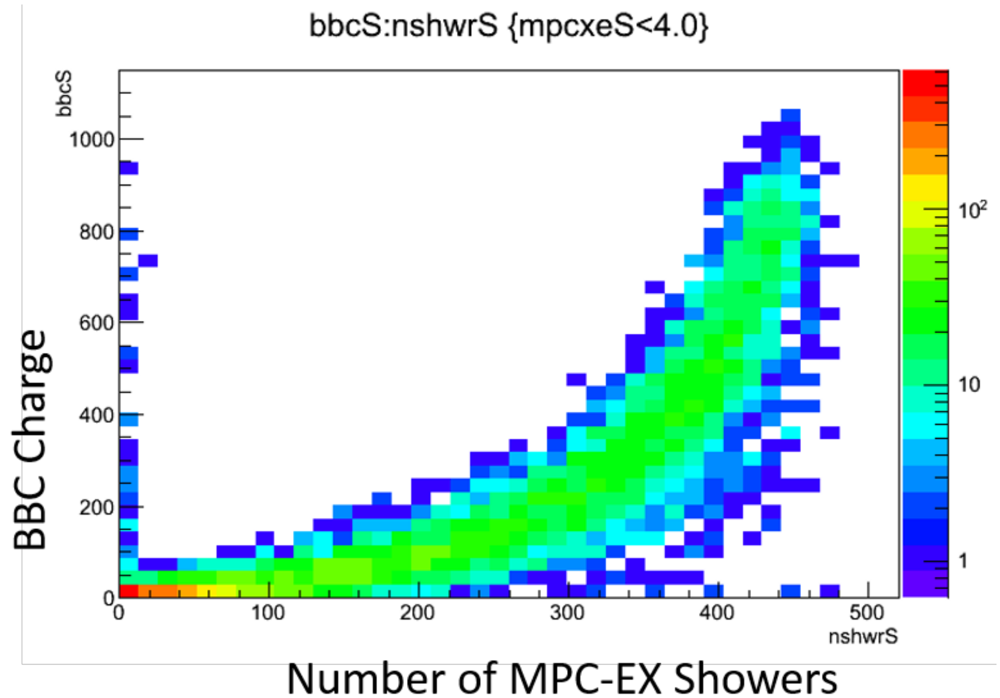

FIGURE 6: CORRELATION OF THE BBC CHARGE WITH THE NUMBER OF SHOWERS FOUND IN THE MPC-EX FOR THE PHENIX SOUTH ARM IN RUN-16. THE LEVEL-1 DELAY FOR THE SOUTH MPC-EX HAS BEEN CORRECTED AND TIMED-IN RELATIVE TO THE BBC, PRODUCING A CLEAR CORRELATION BETWEEN THE NUMBER OF SHOWERS AND THE BBC CHARGE.

In addition to the investigations of the Run-15 data, the detector timing exercise in Run-16 Au+Au collisions was completed for both the north and south MPC-EX. The SVX4 serial download was corrected, and following the original plan from Run-15 a Level-1 delay scan was performed around the nominal value of 40BCLKs, as expected from the VTX strip pixels. Due to the higher multiplicity in Au+Au collisions, and the improved detector performance, it was not necessary to look at MIPs in the MPC-EX but the MPC-EX summed energy (or number of showers) could be compared directly to the BBC charge. The appropriate Level-1 delays for the north and the south were quickly identified.

Figure 6 shows the correlation between the BBC charge and the number of MPC-EX showers in Run-16 Au+Au data. This correlation changes dramatically with a variation of the Level-1 delay of +/-1BCLK, clearly indicating the correct beam crossing, as shown in Figure 7.

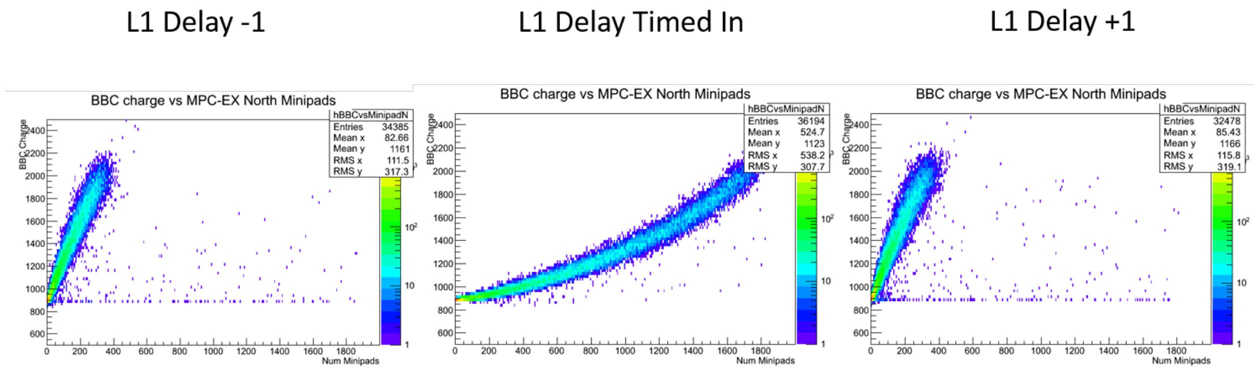

FIGURE 7: CORRELATIONS BETWEEN THE BBC CHARGE AND NUMBER OF HIT MINIPADS IN RUN-16 AU+AU DATA. THE CENTRAL PANEL SHOWS THE CORRELATIONS WITH THE LEVEL-1 DELAY ADJUSTED SO THE BBC AND MPC-EX ARE ON THE SAME BEAM CLOCK, WHILE THE LEFT AND RIGHT PANELS SHOW THE EFFECT OF SHIFTING THE LEVEL-1 DELAY AWAY FROM THE CORRECT VALUE.

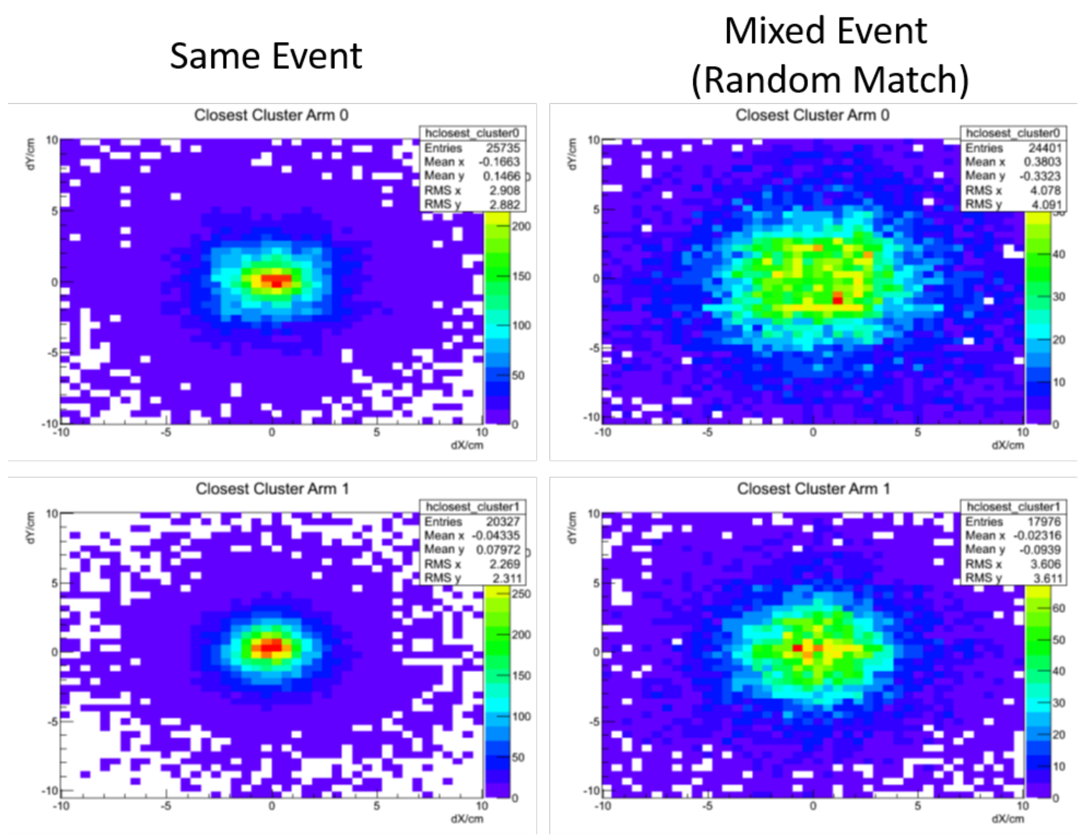

FIGURE 8: A COMPARISON OF THE SPATIAL MATCHING IN X (HORIZONTAL) AND Y (VERTICAL) BETWEEN AN MPC-EX SHOWER AND A CLUSTER IN THE MPC IN RUN-16 DATA. THE COMPARISON IS MADE FOR MPC CLUSTERS WITH AN ENERGY > 1.5 GEV, AND FOR MPC-EX SHOWERS WITH ENERGY IN SEVEN OR MORE LAYERS OF THE MPC-EX. THE LEFT HAND COLUMN SHOWS THE MATCHING IN LOW-MULTIPLICITY AU+AU EVENTS (BBC CHARGE < 200) IN THE NORTH ARM (TOP) AND SOUTH ARM (BOTTOM). THE SAME COMPARISON IS SHOWN FOR MIXED EVENTS WHERE THE MPC-EX SHOWER AND MPC CLUSTER ARE DRAWN FROM DIFFERENT EVENTS. THE SAME-EVENT COMPARISON IS SIMILAR TO THAT EXPECTED FROM SIMULATIONS AS SHOWN IN FIGURE 5.

Finally, Figure 8 shows the spatial matching between the MPC-EX shower axis in Run-16 low-multiplicity Au+Au data for both same-event and mixed event correlations. The same-event correlations compare very favorably with the correlations in simulation shown in Figure 5, and indicate that the MPC-EX shower and MPC cluster are from the same electromagnetic shower.

Conclusions on MPC-EX Status

Based on the results of the study of the Run-15 data, in light of the error discovered in the SVX4 serial download, and the successful timing exercise in Run-16, we are forced to conclude that in the Run-15 data the MPC-EX was mistimed with the rest of PHENIX, and this data is not useful for the MPC-EX physics program. This same study demonstrates that in Run-16 the MPC-EX is fully ready to take advantage of the 200 GeV d+Au run.

Short Discussion on Causes

There are a number of factors that directly contributed to the misconfiguration of the MPC-EX detector and failure to discover the error in a timely fashion in Run-15, including:

- The most immediate cause of the problem was a mis-edited serial command string for the SVX4 readout chips that was kept in a text configuration file and edited by hand. There were no procedures for crosschecks or verifications of configuration changes in place for MPC-EX operations in Run-15.
- The detector readout system and firmware were unavailable until well after the start of Run-15. The detector had been bench-tested with pre-production readout hardware but was not fully tested on the bench with the production readout electronics. This led to a number of issues in the MPC-EX readout being discovered very late, while data taking was underway, and contributed a great deal of confusion in attempting to understand the MPC-EX data as it was coming in. Issues associated with bad/dead/hot channels and errors in the readout led to an erroneous conclusion about correlations with the BBC charge in Run-15 p+p running, which has rather low multiplicity in the MPC-EX.
- Issues with readout noise and the production of the MPC-EX micromodules were discovered very late, almost at the start of the production process. This was due primarily to insufficient testing of the pre-production micromodules. While the issues were successfully addressed and the detector itself has very low noise and a very high fraction of live channels, this delayed production of the micromodules and pushed back the installation of the MPC-EX to the very last minute.
- The original schedule for the MPC-EX at the time of the proposal called for the MPC-EX data collection period to happen over two runs, with the first run starting with Au+Au for detector commissioning and finishing with polarized p+p, and the second run continuing with p+Au. In retrospect this would have been optimal, allowing time for the detector to be tuned up in a higher multiplicity environment. (This was not fully appreciated by the MPC-EX group at the time.) Delays in the start of

funding (six months after the final review), detector construction (described above) and changes to the RHIC running schedule and the future of PHENIX combined to squeeze the MPC-EX data taking period into a single RHIC run with no explicit commissioning time. In retrospect the schedule that resulted from delays on the front end and schedule changes on the back end was not reasonable.

- While there was a detailed plan for commissioning and timing in the detector, the emphasis was placed in knowing the timing of a similar system (the PHENIX VTX) and the use of the MIP peak in the MP-EX raw ADC data. Insufficient attention was paid to correlations between the MPC-EX and other detector systems during the run, such as the BBC and MPC detectors. In particular, the online monitoring in Run-15 focused solely on the MPC-EX detector and detecting issues with the detector caused by the firmware issues. Correlations between the MPC-EX and the BBC and MPC were relegated to an offline analysis of the data. In Run-16, the online monitoring has been redone to include correlations with the MPC and BBC so that these are monitored in real-time.
- The offline framework to read in the raw data and present it to users in a manageable format was organized in a way that abstracted the detector from the actual hardware. This abstraction made it very complicated to use as an aid in commissioning the detector in its imperfect state in Run-15. After Run-15, the MPC-EX group came together to develop a new framework to address these issues, and this framework is in place for Run-16.

Many of the issues and lessons learned described above center on schedule and time. An extremely tight schedule and late delivery of key components created a situation that required the MPC-EX group to work furiously to meet the deadlines required to get the detector working in Run-15, leaving very little time for key personnel to think carefully through each step of the process rather than move on to the latest crisis. In many ways this was not a failure of the detector hardware itself, but a failure to leave adequate time to allow the human element of the project to commission and understand a new detector.

While additional manpower, especially at the R&D and prototyping stage, might have lessened some of the schedule issues with the MPC-EX, we do not believe that additional manpower alone would have completely prevented the failure. Without a commissioning run for a new detector the potential will always exist for the detector operators to fail to configure the detector optimally for immediate physics data taking.

Glossary of Acronyms

ADC	Analog-to-Digital Converter
BBC	Beam-Beam Counter, forward charged-particle detector in PHENIX
BCLK	Bunch CLocK tick, a nominally 106.5 ns separation, synchronized to timing of beam bunches in RHIC
CMM	Coordinate Measuring Machine
GL1	Global Level 1, term for the primary trigger system of PHENIX
HFT	Heavy Flavor Tracker, four-layer silicon tracking detector in STAR
HLT	High Level Trigger, system for real-time track reconstruction in STAR
MIP	Minimum-Ionizing Particle
MPC	Muon Piston Calorimeter, a PHENIX Lead-Tungstate calorimeter in the forward direction
MPC-EX	Muon Piston Calorimeter EXtension, a Tungsten-Silicon preshower located in front of the PHENIX Muon Piston Calorimeter
PXL	two innermost silicon PiXeL layers of the STAR HFT
RDO	ReaDOut board, intermediate electronics board to transfer data from the front-end electronics into the final data stream
SVX4	Front-end readout chip of the MPC-EX
TPC	Time Projection Chamber, a tracking detector in STAR
VTX	PHENIX silicon-based tracking detector for VerTeXing